

AD-A062 373


DEFENCE AND CIVIL INST OF ENVIRONMENTAL MEDICINE DOW--ETC F/G 14/5
GUIDED TO LANDSAT IMAGE ENHANCEMENT PROGRAMS ON THE CCRS DECSYS--ETC(U)
1978 M M TAYLOR, S SCHLIEN

UNCLASSIFIED

DCIEM-TR-78X22

NL


1 OF 1
AD 62373



END
DATE
FILMED

3--79

DOC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A



ADA062373

DDC FILE COPY

6 GUIDE TO LANDSAT IMAGE ENHANCEMENT PROGRAMS
ON THE CCRS DECSYSTEM-10/MAD DISPLAY*

10 M.M./Taylor S./Schlien

Defence and Civil Institute of Environmental Medicine
Box 2000, Downsview, Ontario M3M 3B9

and

S.Schlien

Canada Centre for Remote Sensing
2464 Sheffield Road
Ottawa, Ontario, K1A 0Y7

14 DCIEM-TR-78X22

9 * This is DCIEM Technical Report, Number 78-X-22

12 59p.

11 1978

78 12 13 054
✓ 406 986

See

CHAPTER 1

INTRODUCTION

This report describes a suite of programs for colour enhancement of Landsat images. The programs were developed over the period 1973-77 cooperatively between DCIEM and CCRS for use on the CCRS DECsystem-10/MAD Display.

The report is both a user's guide and an aid for programmers who may wish to understand, alter, or transport the routines. By following the guide, a user should be able to take any Landsat image on a Computer Compatible Tape (CCT) and produce on the Multispectral Analyser Display (MAD) display screen an enhanced image suited to his own individual requirements. The enhancement uses a "principal components" transformation of the Landsat spectra, determined by the statistics of an interesting area, and followed by another transformation which reduces the data to colour values. This report is intended to aid users in handling the routines which implement the enhancement, not as a scientific or technical description of the underlying theories.

In the initial stages of the project, it was thought that a single principal components transformation applicable to almost any Landsat image could be used to reduce to three major components the information in the four spectral bands. A simple mapping of the transformed data into the three-dimensional space of colour would then produce an enhanced display. This concept proved impractical, but some of the earlier image enhancements allowed one to see many new structures which were barely visible in the original three-colour Landsat composite.

Experience with the early version soon demonstrated that no single transformation could provide strong enhancement for all scenes. An enhancement based on data from one scene usually would provide unsatisfactory results on another scene, because of differences in illumination, atmospheric haze, type of terrain, and so forth. In addition, it became clear that there were at least two different types of enhancement: one type in which the general content of a scene was enhanced, and one in which a particular characteristic ground cover type, such as softwood or water, was enhanced without consideration of the rest of the scene. Even though there was and remains some possibility of providing a general enhancement for scenes having similar illumination and atmospheric haze, it seems unlikely that any simple enhancement procedure will give a common enhancement for the specific terrain types of interest to different users. Accordingly, the development of the package has concentrated on two major lines: correction and preprocessing of the imagery data, and improved versatility in the production of enhancements to the taste of the users.

Today the software offers many additional capabilities. There are provisions for (1) correction of sensor miscalibration, (2) haze removal, (3) sun angle compensation, (4) geometric correction, (5) spatial filtering and (6) training for a supervised classifier. The user can generate hard copy photographs of any image that has been displayed on the Multispectral Analyser Display (MAD) using either the Electron Beam Image Recorder (EBIR)

Enhancement Guide: Introduction

or the Continuous Scan Film Recorder (CSFR). The programs give the user many options and allow him to create his own principal component transformations and to control colour contrast and image sharpening. The programs have been used for many applications, such as mapping forests to show areas of fire hazard, ice reconnaissance, mapping diseased forest regions, analyzing turbidity structures in rivers, or just determining what information is available in a Landsat image.

The report serves three purposes; (1) it is a user's guide on how to run the programs; (2) it describes the algorithms used by the program and (3) it gives a general overview of the program software structure. The first half of the report describes the programs from the viewpoint of the user, and the second half describes them as seen by the programmer.

The programs are mostly written in heavily commented FORTRAN IV, and some of the software has been transferred successfully to other computer installations. The package routines will, as a rule, be made available on request to any Canadian non-profit, educational, or Government organization. Arrangements for use by commercial organizations should be made through CCRS or DCIEM.

1.1. GETTING STARTED

In addition to having a valid project-programmer account and a Landsat computer compatible tape, the user should acquire a little familiarity with the MICA system. The MICA system is a collection of programs for thematic mapping. They have been documented in a series of HELP files stored in the DECSYSTEM-10 HLP: account. To obtain access to these files, type HELP MICA.

The enhancement programs read the Landsat digital image from a disk file which is in a special format on the user's private disk pack. This file may be created using either LOADPK for an old format tape, or CCTLOD for a new format tape.

In addition to the image data file, the user must have in the account another file called CONPAR, which is created by running the system program SETPAR. (Type HELP SETPAR for instructions on running SETPAR). Not all the parameters in the CONPAR file need be set. The necessary ones are MSSNA (the name of the image data file), MSSTR (the name of the private pack containing the image data file), and MSFOR (indicates the format of the image data file). If these parameters have been set, and the image data file loaded onto the private pack, the enhancement programs may be run.

Section 1	
b. H. Section	
<input type="checkbox"/>	
<input type="checkbox"/>	
Section 2	
Section 3	
Section 4	
Section 5	
Section 6	
Section 7	
Section 8	
Section 9	
Section 10	
Section 11	
Section 12	
Section 13	
Section 14	
Section 15	
Section 16	
Section 17	
Section 18	
Section 19	
Section 20	
Section 21	
Section 22	
Section 23	
Section 24	
Section 25	
Section 26	
Section 27	
Section 28	
Section 29	
Section 30	
Section 31	
Section 32	
Section 33	
Section 34	
Section 35	
Section 36	
Section 37	
Section 38	
Section 39	
Section 40	
Section 41	
Section 42	
Section 43	
Section 44	
Section 45	
Section 46	
Section 47	
Section 48	
Section 49	
Section 50	
Section 51	
Section 52	
Section 53	
Section 54	
Section 55	
Section 56	
Section 57	
Section 58	
Section 59	
Section 60	
Section 61	
Section 62	
Section 63	
Section 64	
Section 65	
Section 66	
Section 67	
Section 68	
Section 69	
Section 70	
Section 71	
Section 72	
Section 73	
Section 74	
Section 75	
Section 76	
Section 77	
Section 78	
Section 79	
Section 80	
Section 81	
Section 82	
Section 83	
Section 84	
Section 85	
Section 86	
Section 87	
Section 88	
Section 89	
Section 90	
Section 91	
Section 92	
Section 93	
Section 94	
Section 95	
Section 96	
Section 97	
Section 98	
Section 99	
Section 100	

1.2. OVERALL STRUCTURE OF THE ENHANCEMENT PROGRAMS

The enhancement package consists of three major programs, ENHANC, EIGENC, and PIXSET-PIXROL.

ENHANC performs sensor correction and other corrections to the radiometric data;

EIGENC computes the principal components transformation;

PIXSET-PIXROL displays the picture.

It follows that the user must run the programs in order. In normal practice, ENHANC automatically calls either EIGENC or PIXSET-PIXROL, and EIGENC automatically calls PIXSET-PIXROL, so that the user need not be concerned with running the individual programs in the package. These programs pass information to each other through intermediate disk files named, *.CFL, SUN.FAC, BANVEC.LIM, XFILE.NAM, *.XFM, and PX.COM (where * denotes any name up to 6 letters chosen by the user).

The ENHANC program creates a lookup table for sensor calibration, linearization, haze correction, and sun angle compensation. The table is stored in the file *.CFL and the sun angle correction as a single number in the file SUN.FAC in the user's account. The first 6 letters of the *.CFL file are always set to the name of the video file which the user is working with and whose name was initially set in MSSNA. For new format JSC (Johnson Space Center format) tapes, which have already been sensor calibrated, linearized, and destriped, these lookup tables are not computed and ENHANC does nothing but point the user to the other programs. Haze and sun angle correction is therefore not available for images taken from JSC tapes.

The EIGENC program finds the principal components transformation for the user's image and stores it in a file with a .XFM extension. The transformation is estimated from the statistical properties of the image. EIGENC uses the file BANVEC.LIM, which contains limiting values for the intensities acceptable in the four sensor bands. Data falling outside these limits are not included in the statistics, so that the limits can be used for making special-purpose enhancements directed at specific terrain types. The function of the BANVEC.LIM file is discussed in more detail in a later section.

The programs PIXSET and PIXROL read the transformation file, the calibration lookup tables, enhance the image and display it on the MAD under the user's control. PIXSET and PIXROL were originally one program, but were later split into two programs to reduce the computer memory requirements. The user controls the program by running PIXSET which communicates with PIXROL through an intermediate file PX.COM.

CHAPTER 1

INSTRUCTIONS FOR USING THE ENHANCEMENT ROUTINE PACKAGE

1.1. INTRODUCTION

ENHANCE is the first routine of the package. Since each routine calls the next one without user intervention, the user may consider the sequence of routines as a single "enhancement routine". The actual routine called first is really the calibration routine.

The user enters the suite of programs by typing

ENHANCE

The computer responds:

HAVE YOU CHANGED FRAMES SINCE LAST RUNNING ENHANCE?

The user types either Y or N depending on whether a new calibration is required. Usually, one calibration suffices for any one frame, but fails when applied to a different frame, because of camera drift and differences in film angle and development time.

WHEN DONE, DO YOU WANT TO CALCULATE STATISTICS (S), OR TO DISPLAY AN IMAGE (D)?

If the user is running the program for the first time on a particular frame, he must type S. This will instruct ENHANCE to call EIGENC when it is finished, so that a *XEM file can be made for the display. If D is typed, ENHANCE will call FIRST-PIECE when it is finished, bypassing the computation of the principal component transformation matrix stored in *XEM. It is usually not advisable to use an old *XEM file with a new calibration, but sometimes this is done when special circumstances (e.g., translation or enhancement to a new image) are involved.

IS THIS A NEW FORMAT CUT?

If the type is not a standard 45C format tape but a new (old format) tape, the user should type N, and the program will start going through the normal calibration procedure. Otherwise the program goes directly to EIGENC or FIRST-PIECE depending on whether the user typed S or D in the previous question.

We shall assume that the user is using an old format tape type and describe the calibration procedure followed by the program.

The next set of questions is used for computation of rotating the film angle, so that the image can be normalized to the horizontal. This would be found as a latitude of 50 degrees at the equator (about 75 degrees). Normalization is done by multiplying all intensity values by a constant after they have been linearized. The constant is the ratio of the ratio of the sum angle

CHAPTER 2INSTRUCTIONS FOR USING THE ENHANCEMENT ROUTINE PACKAGE2.1. ENHANC

ENHANC is the first routine of the package. Since each routine calls the next one without user intervention, the user may consider the command R ENHANC as if it called an "enhancement routine". The actual routine called first is really the radiometric calibration routine.

The user enters the suite of programs by typing

R ENHANC

The computer responds:

HAVE YOU CHANGED FRAMES SINCE LAST RUNNING ENHANC ?

The user types either Y or N depending on whether a new calibration is required. Usually, one calibration suffices for any one frame, but fails when applied to a different frame, because of sensor drift and differences in sun angle and atmospheric haze.

WHEN DONE, DO YOU WANT TO CALCULATE STATISTICS (S) OR TO DISPLAY AN IMAGE (D).

If the user is running the programs for the first time on a particular frame, he must type S. This will instruct ENHANC to call EIGENC when it is finished, so that a *.XFM file can be made for the display. If D is typed, ENHANC will call PIXSET-PIXROL when it is finished, bypassing the construction of the principal components transformation matrix stored in *.XFM. It is usually not advisable to use an old *.XFM file with a new calibration, but sometimes this is done under special circumstances (e.g. transferring an enhancement to a new image).

IS THIS A NEW FORMAT CCT?

If the tape is not a standard JSC format tape but a raw (old format) tape, the user should type N, and the program will start going through its sensor calibration procedure. Otherwise the program goes directly to EIGENC or PIXSET-PIXROL depending on whether the user typed S or D to the previous question.

We shall assume that the user is using an old format raw tape and describe the calibration procedure followed by the program.

The next set of questions is used for computing or entering the sun angle, so that the image can be normalized to the sun angle that would be found at a latitude of 50 degrees at the equinox (about 37 degrees). Normalization is done by multiplying all intensity values by a constant after they have been linearized. The constant is the ratio of the sines of the sun angle

and the normalization sun angle. The sun angle for the image is determined from the day and latitude which can be calculated from the Fleming Centre Numbers. These are three numbers describing the track (TT), frame number along the track (FF) and 18-day cycle number (CC). Together, they uniquely specify an image. The sun angle correction is not exact, but may be in error by 1 or 2 degrees under some circumstances. If the user knows an accurate sun angle, it can be entered directly instead of being computed.

DO YOU KNOW THE SUN ANGLE (Y-N)?

If the user enters Y, the program asks for the actual angle. Otherwise it asks for the data with which it can compute the angle.

Either:

ENTER SUN ANGLE IN DEGREES AND DECIMAL DEGREES

or:

LANDSAT -1 or 2 ?

The user specifies whether his frame is from Landsat 1 or 2.

ENTER FLEMING CENTRE FOR FRAME TT,FF,CC

The sun angle correction is determined and copied into the file SUN.FAC .

TYPE T FOR TAPE , D FOR DISK

ENHANC is the only program of the suite which accepts Landsat data on tape (old and new formats). If the user specifies T for tape, the program asks

WHAT NAME IS YOUR TAPE ASSIGNED AS ?

The program must know the logical name assigned to your tape drive so it can open a channel to the appropriate device. For example, if the user mounted his tape as follows:

MOUNT TCA:TAPE/VID:RS0000

then the tape drive is assigned logical name TAPE.

More likely, the user's input data is on disk. The program reads the user's CONPAR file and extracts the file name, disk pack name, and format of the file. The program then opens a channel to this file.

PICTURE BOUNDS FOR RADIOMETRIC CALIBRATION

TYPE STARTING LINE AND PIXEL

The calibration program computes a look up table from the statistics of the sensor responses. It is not necessary to sample all the pixels in the image; the program works reasonably well with 100,000 samples provided they are fairly typical of the image. The program samples from the image a rectangular area whose top left hand coordinates are specified by the starting

line and starting pixel. (Lines on the disk file always number 1 to N starting from the top line in the stored image. Pixel numbers go from 1 to 3200 starting from the left edge.)

TYPE NUMBER OF LINES AND OF PIXELS

The user specifies the dimensions of the sampling area. (100 lines by 3000 pixels is usually adequate provided that if the enhancement or the display is going to incorporate water, then the calibration area must include a reasonable amount of water.)

DO YOU WANT AUTOMATIC HAZE CORRECTION?

The user can minimize the haze effects by typing Y. The haze removal is only approximate, and if one is working with a single image, it may be omitted. Haze removal helps achieve common conditions across several frames, so that if one intends to use a single enhancement transformation across several images, then haze removal is recommended. Within a single image, however, the data are often distributed in such a way that the enhancement is better without haze removal. The user is advised to try it both ways in the early stages, so as to get a feeling for the kinds of effects obtained with and without haze removal.

The haze corrections are determined by finding the lowest intensity level in each sensor band, in the sampling area specified by the user. The program works on the assumption that somewhere in the scene there is a dark or low-reflectance region in at least one of the bands. The minimum values in the four bands are compared to each other and the one that is lowest compared to its expectation is chosen as an indicator of haze level. The four bands are then corrected by reference to this chosen band so that in the chosen band the lowest value is given a zero level in the lookup table, while the lowest values in the other bands will be zero or positive.

The program types out

CALIBRATING SCENE

APPROX 10 CPU SECONDS ...

and proceeds to accumulate statistics. The estimated time depends on the number of lines and pixels to be sampled.

Once the lookup tables are computed, they are written on a *.CFL file whose name is the same as the user's image file and the program goes to either EIGENC or PIXSET. A description on how the sensor calibration is performed is given in the next section.

2.1.1. SENSOR CALIBRATION

In each of the four bands (4,5,6, and 7) the Landsat scene is viewed by six individual sensors which together scan one swath of six lines. Because the sensor responses differ slightly from one another, the resulting image has horizontal striding which is periodic every six lines. The calibration program determines the sensor responses relative to each other and produces a lookup table which equalizes them. The same lookup table is used to incorporate other sensor corrections. These are decompression (linearization) from the quasi-logarithmic compression employed in transmitting the Band 4, 5, and 6 data to Earth, an approximate correction for seasonally varying sun angle, and an approximate haze correction. Both the latter two corrections are optional.

The equalization is accomplished by computing the intensity histograms for each of the six sensors and for each of the four bands. Since the six sensors view almost the same area, the histograms should be practically identical. However, due to the different responses of the sensors, some of the histograms are shifted or are broader than normal. The sensor equalization scheme determines a transformation which makes the statistics from each of the six sensors look the same.

Once the histograms are computed, they are cumulated, and normalized to a distribution function. The distribution function for the second sensor in each band is arbitrarily taken as the standard, and the intensities of the other sensors are modified so that the functions completely match. This scheme departs from the usual CCRS calibration procedure in the sense that it is nonparametric and does not assume a linear relationship between the sensors. At the end of this process, it remains to incorporate the linearization transformation, sun angle correction, and the optional haze correction.

2.2. EIGENC

The intensities in the four spectral bands from Landsat are highly correlated with each other. In other words, the images derived from band 4 and band 5 look very much alike, as do the images derived from band 6 and band 7. There is also a good deal of similarity, though not as much, between the images in the visible bands (4 and 5) and those in the infrared bands (6 and 7). If the bands are used individually in preparing a coloured display, only three of the four can be presented in one picture. The remaining band must be viewed separately. In order to be able to display all the information that is available in the Landsat image it is desirable to eliminate the redundant information and to incorporate the essential information in three "channels" for display in colour.

Various simple transformations of the data, such as taking the ratios among the bands, or taking sums and differences, have been used with dramatic effect, but they are not usually optimum for any particular image. It has been shown mathematically that under Gaussian assumptions an optimum transformation can be found. It is commonly known as the principal components transformation. This transformation produces a new set of features or vectors which are linear combinations of the original Landsat bands. The data described in terms of these vectors are uncorrelated and most of the information is contained in the first one or two components. If the user generates black and white images from each of these vectors, each image contains unique information which is not present in the other images. Furthermore, the images of the last vectors seem more uniform and very noisy.

Unfortunately, there does not seem to be any one principal components transformation which is ideal for every Landsat image and every user. The transformation depends entirely on the data used to determine the correlations among the bands. The details of the correlations change from scene to scene and from place to place within a scene. If the user wishes to enhance the variations within a forest, it will not be useful to gather data from a large lake, or from cloud cover. The program EIGENC has therefore been designed to gather the statistics from a user-selected portion of Landsat data, and to determine the principal components based on these statistics. The resulting transformation is recorded in a *.XFM file which the user will retrieve when he runs the program PIXSET. The name given to this *.XFM file is chosen by the user, and PIXSET will automatically use the given name when called by EIGENC, unless told otherwise by the user (see PIXSET command "XFM" below).

The user enters the EIGENC program either automatically from ENHANC or by typing R EIGENC. The program proceeds with the following dialogue.

DO YOU WANT TO USE ALREADY CUMULATED
DATA TO MAKE A TRANSFORMATION TABLE?

If the user is running the program for the first time on a particular scene, the answer must be N. Using already cumulated data means combining

covariance matrices that have been produced earlier by of EIGENC, whether on this image or on others. The user has the option of combining these, and if he has typed Y, the program asks:

NAME OF COVARIANCE FILES (C.R. TO END LIST)

*

and accumulates the statistics from the *.XFM files which the user designates. (These files were previously produced by running EIGENC).

Assuming that the user is not working with already cumulated data, then the program proceeds to read the user's CONPAR file to find out the name of the Landsat data file, the private pack containing this file, and the mode of reading the file. (The CONPAR file was created and edited by running the program SETPAR.) If the video file is being read in new format mode the program also types

READING NEW FORMAT IN 7 BIT PRECISION

(In the case of new format JSC tapes, the user can read the data from disk in either 6, 7, or 8 bit precision, by using SETPAR to set the MSFOR parameter. As long as the data is read consistently in the same manner by both EIGENC and PIXSET-PIXROL, the user will obtain satisfactory results. Very strange pictures result from an inconsistent use of MSFOR, such as using one precision when calculating the principal components with EIGENC and another when displaying the picture with PIXSET-PIXROL. It is probably best to read the data in full precision, 8 bit mode, by setting the MSFOR parameter to 4 when running SETPAR. The intensity levels in a standard JSC tape have been scaled to 256 grey levels.)

In the following section, typical user responses have been used, to indicate the appropriate formats.

The program next asks:

LIMITS ON ACCEPTABLE DATA

A=ACCEPT ALL

D=DEFAULT (LAST LIMITS DESCRIBED)

N=NEW

The user has the option of ignoring pixels whose intensity level in any one of the four bands does not fall within a specified range. This is useful when the user is creating a transformation to enhance a specific ground cover type such as water. To do this, the user sets the limits so that only water pixels are acceptable. Setting the upper limit to 6 or 7 on Band 7 usually accomplishes this result. If the user wishes to set certain limits, then he types N, meaning "new". The program will then ask for the lower and upper limits. In the following example, we assume that the user wants to enhance everything except water and clouds, so he sets the lower limit on Band 7 to 7 (excluding water) and the upper limit on Band 4 to 40 (excluding clouds):

LOWER BOUNDS FOR 4 BANDS, IN ORDER

0 0 0 7

and then

UPPER BOUNDS FOR 4 BANDS IN ORDER

40 63 63 63

The limits are written on the file BANVEC.LIM, which will be used in future whenever the user types D for default. The same limits will later be available to the PIXSET-PIXROL display routines. The program writes out the limits on the user's terminal

LOWER LIMITS = 0 0 0 7

UPPER LIMITS = 40 63 63 63

ARE THESE OK? Y

If the user types Y, the program carries on, otherwise it asks again for the limits.

TYPE NAME OF COVARIANCE FILE, DEFAULT IS TNFRB.

FILNAM

The user may give his covariance file any name, but the program will always give it a .XFM extension. If the user just hits carriage return, the program will choose the default name TNFRB.XFM. If the user specifies a file which already exists, the program will type:

FILNAM.XFM EXISTS ALREADY,

DO YOU WANT TO ADD TO IT (A)

DELETE IT BEFORE REWRITING (D)

OR USE A NEW NAME (N)

D

The program will either read in the accumulated statistics in that file, reset these statistics to zero, or ask the user to enter the new file name, depending on whether the user types A, D, or N.

The program now attempts to open a channel to the Landsat video file which is stored on the user's private pack.

At this point, the user now has to specify the sampling area from which the program gathers the appropriate statistics. The user should be sufficiently familiar with his image that he can specify this area in terms of line and pixel coordinates on the disk file. The area which the user selects affects the appearance of the resulting enhancement. For example, if the sampling area has no water bodies, then the water bodies in the

enhanced image will tend to take saturated colours or be black. If the user is very interested in mapping different types of snow, the sampling area should include as much snow as possible. The user will specify one or several rectangular areas from which to gather the statistics. The area is defined in the following manner.

NUMBER OF PIXELS ACROSS, NUMBER OF LINES DOWN = 200,200

The actual number of lines or pixels that will be read from the disk file are specified here. The number of pixels per line must be no larger than 558. Note that the program has the capability of sampling every second, third, ... pixel or line, which effectively expands the sampling area. Typical numbers selected are 200 by 200.

STARTING LINE, STARTING PIXEL = 1,1200

The user specifies the top left hand corner of the sampling area to be the 1200th pixel on the first line of the image.

LINE DECIMATION M OUT OF N = 1,1

PIXEL DECIMATION I OUT OF J = 2,5

DO YOU WANT LOW PASS FILTER ? Y

The "low pass filter" averages a 2X2 pixel region instead of taking single pixels. It is valuable in two situations. Where there is a considerable amount of quantization noise, as might be the case in a clear water area, averaging reduces this noise and increases the inter-band correlation. Where the terrain itself has considerable variation from pixel to pixel, as do some forest areas, averaging has the effect of emphasising the lower frequency components of the spatial covariance, and thus gives the final enhanced display a smoother, more coherent appearance. In very few circumstances will it be beneficial not to use the low pass filter, since if there is neither quantization noise nor rapid terrain fluctuation, the covariances will be little changed by the use of the filter. One situation where low pass filtering should not be used is if there exist in the image many narrow linear features which are of interest to the user. Low-pass filtering would eliminate these from the statistics. Urban regions provide an example. The routine runs more slowly if the low-pass filter is used.

LOG OR LIN TRANSFORMATION ? LOG

The standard enhancement uses a logarithmic transformation of the data, since in most cases the data in logarithmic form conform more closely to the Gaussian assumption. There may be cases, however, in which the linear form gives a better enhancement. These cases can only be discovered by trial. If in doubt, answer LOG. Whichever transformation is used in EIGENC, the same transformation MUST be used when displaying an enhanced image using PIXSET-PIXROL. Normally, this is taken care of automatically, but since the user can alter the transformation used by PIXSET-PIXROL, he should be aware of the potential problem.

The program proceeds to accumulate the means, covariances, compute the principal components transformation and write the results in the user's *.XFM file. This procedure usually takes a few minutes. It then asks:

DO YOU WANT TO DO MORE WITH EIGEN?

If the user answers "Y", the program goes back to the beginning, allowing the user to include the new *.XFM file in a cumulated statistics file, or to make another *.XFM file for a new area or with new limiting values. Otherwise, it calls PIXSET to display a picture.

Assuming the user has produced various covariance files (*.XFM) in his account, he has the opportunity to mix the statistics from several of them and produce a new *.XFM file. Some good enhancements have been obtained by 'mixing a bit of Vancouver' with an existing transformation. There is, however, no set procedure for obtaining specific results. This option was originally introduced in order to find some master transformation which would be adequate for all Landsat images.

One circumstance in which combining statistics is useful is in the case that a particular terrain type of interest occupies only a small portion of a larger image which must all be enhanced. A town is a common example. It is often helpful to make one *.XFM file for the larger area, and one or more *.XFM files for the interesting terrain sampled quite densely and possibly repeatedly using several passes. The several *.XFM files can then be combined, using the option YES when EIGENC asks whether you wish to use cumulated data.

LOG ON NEW TRANSFORMATION ? LOG

The standard enhancement uses a logarithmic transformation of the data. Since in most cases the data in logarithmic form conform more closely to the Gaussian assumption, this may be better, however, in which the linear form gives a better enhancement. These cases can only be discovered by trial. If in doubt, answer NO. Whenever transformation is used in EIGENC, the same transformation MUST be used when displaying an enhanced image using PIXSET-PIXNOL. Normally, this is taken care of automatically, but since the user can alter the transformation used by PIXSET-PIXNOL, he should be aware of the potential problem.

2.2.1. PRINCIPAL COMPONENTS DETERMINATION

A non-mathematical reader may wish to skip the following section, in which some of the details of EIGENC are considered.

The principal components transformation is computed from the mean vector and the matrix of covariances among the 4 band intensities sampled by the program. The mean vector refers to the average brightness of the sampling area, the covariance matrix indicates the amount of variability in that area (diagonal components), and the degree of correlation between the bands (off-diagonal components). These statistics are computed from training samples which the user selects (just as in supervised maximum likelihood classification).

Once the covariance matrix has been obtained, its eigenvalues and eigenvectors are computed using subroutines from a standard matrix package. Since the covariance matrix of real data (as opposed to complex or imaginary) is positive semidefinite, the eigenvalues are always greater than or equal to zero.

A new black and white image can be constructed using any one of the eigenvectors. The eigenvector is treated as a set of weights, one for each sensor band. When the sensor values for a pixel (corrected by use of the lookup table) are multiplied by these weights and summed along with a correction for the mean, a pixel of the new image is created. Colloquially, the new images formed by using the eigenvectors in this way are called "vector images".

The computed eigenvalues reflect the amount of information in each of the vector images. In other words, the image produced by the eigenvector with the smallest eigenvalue appears very grey and lacking in spatial structure. The matrix subroutines conveniently provide the eigenvectors in order of decreasing eigenvalue.

To get the principal components, the eigenvectors are next divided by the square roots of the eigenvalues. This ensures that the variance of a picture derived from any of the eigenvectors is unity. The next step is to determine the mean values of the principal components so that the means can be corrected to a consistent grey level when the vector images are displayed using PIXSET-PIXROL. The last step involves checking the polarity of the vectors, and reversing some of them if necessary. This helps in making more standard enhancements, and avoids situations where lakes and rivers appear bright saturated red colours, and clouds appear dark green or blue. The convention used is such that pixels which are brighter in the original will be brighter in the enhanced colour picture, and that pixels relatively brighter in the infrared will be shown more green.

2.3. PIXSET-PIXROL

The Landsat data is enhanced and displayed using the pair of programs PIXSET-PIXROL. When the user enters PIXSET, either automatically from ENHANC or EIGENC, or by typing R PIXSET, the program types

TYPE HELP FOR CURRENT COMMAND LIST

LAST UPDATED 22 FEB 1976.

!

PIXSET is a command driven program and has many similarities to the program MADCON (see MAD USERS MANUAL, or type HELP MADCON) from which it developed. The exclamation mark indicates that the program is ready to accept any of the 32 commands. Using these commands PIXSET modifies a table which instructs PIXROL how to enhance and display the image. The commands PIXEL, LINE, and MAG allow the user to specify what part of the Landsat image to display. When the user types ROLL, the table is written on a file called PX.COM and the program PIXROL is fetched. PIXROL reads this file and starts displaying the enhanced image according to these parameters. Whenever the user hits the spacebar or types any character on his terminal, PIXROL immediately updates the PX.COM file and returns to the program PIXSET. PIXSET is a communicator passing information from the user to PIXROL. The user can see the list of PIXSET commands by typing HELP. The program will type the following list:

HELP	THIS LIST		
LINE	STARTING LINE NUMBER	PIXEL	STARTING PIXEL NUMBER
LINSEL	LINE SELECTION	PIXSEL	PIXEL SELECTION
MAG	MAGNIFICATION FACTOR	ROLL	ROLL UP SCREEN
NORML	NORMAL DISPLAY	STATU	STATUS
ENHAN	ENHANCEMENT	WRITE	WRITE DISPLAY COMPATIBLE FILE
EXIT	EXIT TO MONITOR	GEOM	SET SCALE TO GEOMETRIC SQUARE
NOGEO	USE DECIMATION	LIMIT	SET NEW BAND LIMITS
BAND	DISPLAY ONE BAND B/W VECTOR		DISPLAY SELECTED VECTOR B/W
TEST	DISPLAY TEST PATTERN	XFM	SELECT TRANSFORM TABLE
SMOOTH	SMOOTH DISPLAY	NOSMOO	DO NOT SMOOTH DISPLAY
GRID	INSERT GRID PATTERN	NOGRID	DO NOT PUT GRID
GRID?	TYPE OUT GRID SETTING	TRUTH	PREPARE GROUND TRUTH FILE
SHARP	SHARPEN DISPLAY	NOSHARP	DO NOT SHARPEN DISPLAY
AREA	USE AREA FILTER	RESET	USE DEFAULT SETTINGS
DRAMA	SET COLOUR SCALE	GRIDOR	SET GRID ORIGIN
NEW	PRINT CHANGE NEWS		

In typing any of these commands, only 5 characters are significant. Hence, one can get the same result by typing GRIDO or GRIDORIGIN, VECTO or VECTOR-NUMBER.

The new user can manage very well with no more than 5 commands. These are LINE, PIXEL, MAG, WRITE, and ROLL. WRITE is described later. The other four are needed to describe and to produce an image of a selected area at a selected size. The following example illustrates the way the user displays

an enhancement of a particular portion of a LANDSAT image. (Lower case letters have been used to distinguish the user's commands from the computer's response, however, the user must use upper case letters for the program commands.)

Enter directly from ENHANC or EIGENC, or

r pixset

TYPE HELP FOR CURRENT COMMAND LIST

LAST UPDATED 22 FEB 1976

!line

STARTING LINE=300

!pixel

STARTING PIXEL=1100

!mag

MAGNIFICATION=0.5

PIXSET allows the user to have positive magnification factors which have any positive value specified to two decimal places. If the magnification is less than one, the program adjusts the scale of the picture by decimation so that the user can view a larger portion of the LANDSAT image. A magnification of 0.17 allows the entire width of the frame to be viewed. Smaller magnifications have no practical value. If the user wishes to zoom into a given area, he sets the magnification factor to a large number such as 3. PIXSET automatically corrects the aspect ratio of the LANDSAT image to give the user a realistically scaled picture. It does not perform other aspects of geometric correction such as earth rotation deskewing, since the data tape may have had these corrections applied, and since they differ from image to image. The geometric correction is simply an expansion by a factor of 1.5 (allowing for both the satellite sampling difference between lines and pixels and the MAD screen ratio) in the vertical dimension. The user may disengage this feature using the NOGEO command described below.

!roll

RUNNING PIXROL

The enhanced image is slowly scrolled up on the MAD. To stop the scrolling and return to PIXSET, the user may hit the spacebar, and the program types:

LINES 300 TO 463 PIXELS 1100 to 2115

RUNNING PIXSET

TYPE HELP FOR CURRENT COMMAND LIST

LAST UPDATED 22 FEB 1976.

!

If the user changed his mind and wishes to continue enhancing the image, he types ROLL. If he wants to check his display parameters, he types STATU.

!status

(note that only the 5 characters "STATU" are significant)

READING VIDEO FILE WINPEG ON WATR
 NEW FORMAT MODE
 NEXT LINE IS 464 FROM PIXEL 1100 TO 2115
 ENHANCED USING MATRIX TNFRB.XFM
 COLOUR SCALE FACTOR=1.
 LOGARITHMIC DATA DISPLAY
 SHARPENING PARAMETER = 1.
 NOT SMOOTHED BY INTERPOLATION
 AREA FILTER DECAY FACTOR =0.8 WEIGHT = 0.3
 TAKING 3 OUT OF 4 LINES AND
 TAKING 1 OUT OF 2 PIXELS
 GRID IS EVERY 100 LINES AND EVERY 100 PIXELS
 STARTING AT LINE 0 AND PIXEL 0.

!exit

To return to the monitor and leave PIXSET, the user types EXIT. When the user restarts the program by typing R PIXSET, the program will read the PX.COM file and come back to the same state, except that any output file on which he may have been writing the image will no longer be open and cannot be appended to.

The enhancements are produced using standard default parameters. Commands are provided for changing some of these parameters, but others were defined either in EIGENC or ENHANC, and cannot be changed at this stage.

The program automatically places horizontal and vertical black grid lines every 100 pixels and every 100 lines. These lines are of great value when the user needs to estimate the line and pixel coordinates of a certain part of the picture. If the lines are a nuisance, or if the user wants to change the grid spacing or origin, he can do this in the following manner.

!grid?
 GRID IS EVERY 100 LINES AND EVERY 100 PIXELS
 STARTING AT LINE 0 AND PIXEL 0.
 !grid
 INSERT GRID LINES EVERY N LINES AND M PIXELS. N,M = 50,50
 !grido GRID ORIGIN LINE, PIXEL = 25,25

Here the user changed the grid spacing to every 50 lines and every 50 pixels, starting at line 25, pixel 25. The grid line substitutes for the named pixel if MAG is 1, and appears at its top and left sides if MAG is greater than 1. For MAG less than 1, the grid line substitutes for one or more pixels starting with the named one.

Note the difference between GRID?, which requests status information on the grid, GRID, which allows the grid spacing to be changed, and GRIDO or GRIDOR (only 5 characters of any command are significant), which allows the grid origin to be set. If the user does not want any grid marks, he types:

!nogrid

The grid marks on the current image on the MAD will not be removed, but no further grid lines will appear on subsequent images the user rolls up. The new values will become defaults as long as the user runs the program in the same account, until he changes or resets them.

Some users who have used MADCON frequently may want direct control of the line and pixel selection parameters. If the user types NOGEO, he disengages the automatic scaling of the image and has the option of setting the line and pixel decimations. An example follows.

```
!pixel
TYPE MAG TO SET SCALE OR NOGEO TO USE DECIMATION
!
```

The user realizes that he made a mistake and proceeds as follows:

```
!ngeo
!pixel
PIXEL DECIMATION I OUT OF J I,J = 1,6
!linel
LINE DECIMATION M OUT OF N M,N = 1,5
!roll
```

The image rolls up with the usual geometric distortion introduced by the nonsquare pixel size (58 by 79 meters).

```
!geom
```

When the user wants to return to the automatic geometric correction mode, he types GEOM.

The user has two modes of displaying LANDSAT image with MAG greater than 1.0. He can either have the pixel elements replicated a certain number of times to obtain the desired scale, or he can have the program linearly interpolate the points between the pixel centres. If the user chooses the first option, the resulting image appears blocky at high magnifications. If he chooses the second option, the image is blurred. Neither result is very appealing, but the only solution is to go to higher resolution data, such as low altitude airborne imagery. The first option is probably useful when the operator is acquiring training samples for supervised classification. The current default is no interpolation, however the user can change this by using the SMOOTH and NOSMOO commands. (See SUPERVISED CLASSIFICATION section.)

```
!smooth
```

```
set up linear interpolation mode
```

```
or
```

```
!nosmo
```

set up replication mode

No interpolation will be done if the magnification factor is 1.0 or less.

The user quickly discovers that the PIXSET and PIXROL programs consume much computer time. When the PDP-10 timesharing system is busy, it may take as much as 15 minutes for a complete image to roll up on the MAD. The user can save his results for fast display by means of the WRITE command. This command is usually, but not necessarily, issued immediately before the ROLL command. The order in which PIXSET commands are used does not matter, but if the user establishes a convention, he is less likely to forget a command he wished to use.

!write

TYPE NAME OF DISPLAY COMPATIBLE FILE. pictur

DO YOU WANT A GREY SCALE?yes

A grey scale is very useful for balancing the colours on the MAD, or for photographic reproduction. The next time the user types ROLL, a copy of the data displayed on the MAD is also written directly on the user's private pack in a file called "pictur". This file is automatically closed after 500 lines have been written. The file can be displayed quickly using the program DSPCIO, or recorded on film by DSPEBR or DSPCFR. (Type HELP DSPCIO (DSPEBR, DSPCFR) for instructions on using these programs.)

If the user knows all the parameters and does not require to see the image, or if the MAD is unavailable, he can store an image for later viewing using the WRITE command. Before entering PIXROL, return to the monitor by typing control-C instead of a command. Then type

as mad nul (of course, using upper case letters)

This command makes all references to the MAD refer to an imaginary device. The programs run happily, but no picture is sent to the MAD.

Special Purpose Enhancements

When the user is only interested in enhancing, say, water bodies, and has tailored his transformation (using EIGENC) for that purpose, then the other ground cover classes may take bright saturated colours. These colours may well make it difficult to see the interesting colour variations within the water. In such circumstances, the user may prefer to have everything else except the water set to a neutral grey colour. This can be achieved using the LIMIT command provided that the user knows which intensity level separates the water from everything else. If the user does not know this, trial and error will quickly determine the correct values. The LIMIT command is exactly the same as the one used in EIGENC, and the default values will be those left there the last time it was used, either in EIGENC or in PIXSET. They are stored in the file BANVEC.LIM (in ASCII, so they could be changed with TECO if desired). If the user wants to see only those areas he sampled in making the transformation matrix with EIGENC, he should type D for default.


```

!limit
LIMITS ON ACCEPTABLE DATA
A=ACCEPT ALL
D=DEFAULT (LAST LIMITS DESCRIBED)
N=NEW
n
UPPER BOUNDS FOR 4 BANDS, IN ORDER
64 64 64 5
LOWER BOUNDS FOR 4 BANDS, IN ORDER
0 0 0 0
LOWER LIMITS = 0 0 0 0
UPPER LIMITS =64 64 64 5
ARE THESE OK? yes
!roll

```

The user assumes that water is never brighter than 5 in band 7. The next time the user rolls up the enhancement, all the pixels with intensity levels greater than 5 in band 7 will appear neutral gray.

The same process can be used in reverse, to determine what areas are to be used to build the statistics for EIGENC. The user can run PIXSET using either an old enhancement transform, or using the NORML command (see below) to obtain a conventional presentation. The limits can be altered by trial and error until just those areas of interest are visible. These limits are now the defaults, and the user can run EIGENC, responding "D" when the limits question occurs. Only the areas of interest will then be used in the statistics. After this pass through EIGENC, PIXSET can be run again to see the special enhancement.

If the user desires, he can view the 'unenhanced' data as it has been read from the disk file.

```

!norml
LOGARITHMIC (LOG) LINEAR (LIN) OR RAW (RAW) DISPLAY?

```

The user has the option of transforming the data to a logarithmic scale. This is the usual default. If the user chooses LIN, the radiance values on a linear scale having a range 10 - 63 will be used. The scale change is required, in order to permit a visible display on the MAD screen without the user having to alter the brightness control on the monitor. If the user chooses RAW, the data are shown as they come from the sensor, with no radiometric correction, swath-banding correction, or other alteration. In most cases, the user will have to alter the brightness setting on the monitor in order to see the picture. From now on, PIXSET will run just like MADCON. The blue green and red guns on the MAD will display the LANDSAT bands 4,5, and 6.

If the user chooses he can display any one of the particular bands in black and white. The radiometric transformation set with the last NORML command will be enforced. If no NORML command has been issued, the transformation will be logarithmic.

!band
TYPE BAND NUMBER 4-7 5

The user can also display any one of the four principal components in black and white.

!vector
ENTER VECTOR NUMBER 1-4 2

The VECTO command may also be used to show colour enhancements using an arbitrary ordering of the principal components, by entering three numbers after the "ENTER VECTOR NUMBER 1-4", as, for example:

!vecto
ENTER VECTOR NUMBER 1-4 1,2,3

This command would produce the conventional enhancement. The first number in the response indicates which vector image is to control the brightness of the enhancement, the second, which controls the red-greenness, and the third, which controls the blue-yellowness. For example, 1,2,4 would substitute the fourth component as controlling the blue-yellow variation. Such unconventional enhancements are used mainly by those who have some understanding of the enhancement technique and want to investigate subtle effects that might lead to new insights leading to better enhancements.

When the user wants to return to the conventional enhancement, he types ENHAN

!enhan

The transformation produced by the EIGENC program is the key factor controlling the resulting enhancement. The user may wish to compare several different transformations that were computed from different samples in a scene. If these transformations have been saved in individual *.XFM files, then the desired file can be selected using the XFM command. For example:

!xfrm
NAME OF TRANSFORMATION test2
!

PIXROL from now on will read the TEST2.XFM file, until the default name is changed. This can be done either with another XFM command, or by running EIGENC, since the name used in the last EIGENC run becomes the default for a PIXSET run.

When the user runs PIXSET for the first time, the default values of the parameters have been set so that some useful enhancement is obtained when he types ROLL. From then on the user changes the parameters, which can either improve or degrade the enhancement. The values of these parameters

will stay as they are set, so long as the PX.COM file on the user's account has not been destroyed.

!reset

Occasionally, a user enters a bad value for one of the parameters, and all the the following images have a disagreeable appearance. The RESET command was introduced for the user who does not remember what has gone wrong and is willing to restore all the parameters to their initial settings. It simply applies all the initial default values, but does not change the name of the *.XFM file or the defaults for the LIMIT command.

2.3.1. COLOUR ENHANCEMENT

The transformation in the *.XFM file condenses the information in the 4 LANDSAT bands into separate vectors. Rather than transmitting these vectors directly to the blue, green, and red guns, PIXROL maps them into a colour space. This mapping adjusts for the human visual response and the CONRAC monitor on the MAD. As described in the paper REMOTE SENSING AND HUMAN VISION (Taylor and Langham, forthcoming), the eye sends visual information to the brain through three independent channels. These channels differ in their information capacity which is determined by their resolution in space, intensity, and time. The brightness channel has by far the best resolution in all three factors, followed by the red/green channel and then the blue/yellow channel. The vector carrying the most information is therefore dispatched to the brightness channel, and the other two vectors modulate the red/greenness and blue/yellowness of the picture.

The colour space fills a three dimensional volume of brightness, hue, and saturation. The user can view a plane of constant brightness on the MAD using the PIXSET command TEST.

```
!test  
ENTER BRIGHTNESS VALUE, 0-63 32
```

A section of the colour space is immediately rolled up onto the MAD. Assuming that the blue, green, red, brightness knobs have been set properly, the user will see a smooth yellow grey blue gradation in the vertical direction and a green grey red gradation in the horizontal direction.

```
!drama
```

The DRAMA command allows the user to increase or decrease the colour contrast in the display. It changes the scale factor in the mapping of the second and third vectors into their red/green and blue/yellow representations. The default setting of 1.0 is probably best for most scenes and observers. Some people, however, tend to prefer more vivid colours, and others prefer more muted scenes. More vivid colours can be obtained with a higher DRAMA value. Values greater than 2.0 are probably not very useful. Values less than 0.5 are useful only if some aspect of the scene that is normally highly saturated is now of interest.

```
!drama  
COLOUR SCALE FACTOR (DRAMA) = 1.3
```

2.3.2. SPATIAL FILTERING

The enhancement programs usually enhance but sometimes obscure thin linear features such as roads and other cultural features which are normally very conspicuous and are excellent reference points. To alleviate this problem, PIXSET has two provisions for applying spatial filtering which will accentuate these features.

The SHARP command was introduced to remove the horizontal pixel overlap. (Recall that pixels have a resolution of 79 meters, but are spaced 58 meters apart in the horizontal direction.) The filter is applied to the four LANDSAT bands before they are transformed to the principal components and the colour space. To change the filter setting the user types

!sharp
SHARPENING VALUE=

A value of 1. corrects for pixel overlap, larger values introduce stronger filtering, a value of zero or less turns the filter off. The filter subtracts a portion of the pixel intensities at the immediate left and right of every pixel. The stronger the filter, the larger this fraction.

$$X'(I) = A * X(I) - B * (X(I-1) + X(I+1))$$

where

$$A = 1 + (\text{sharpening value})/3$$

$$B = (\text{sharpening value})/6$$

$$X(I) = \text{original value of } I\text{th pixel}$$

$$X'(I) = \text{new value of } I\text{th pixel}$$

The filter boosts the high spatial frequencies (i.e. rapid spatial variations) at the expense of the lower spatial frequencies. Since the default sharpening parameter of unity provides good compensation for horizontal pixel overlap, the pixel values are usually more valid with sharpening than without. This is especially true in areas with rapid variation, such as cities or forests. If the sharpening value is set too high, the vertical boundaries and edges will stand out but the rest of the image may look uninteresting. It is unlikely that the user will improve the enhancement with strong sharpening values unless he is only interested in studying vertical lineaments. The user can turn off the filter with the NOSHARP command. The NOSHARP command should ordinarily be used only where the user, for some reason, wants to see pixel by pixel what was returned by the sensors rather than what is in the terrain.

Note carefully the different characteristics of SHARP and NOSHARP, as compared with SMOOTH and NOSMOO. They are not opposites, and can usefully appear in any combination. SMOO and NOSMOO deal with the interpolation between pixel centres needed when magnification is used and either a pixel value must be replicated several times or a series of interpolated values must be used. SHARP and NOSHARP, on the other hand, deal with effects inherent in the data, caused by the design of the satellite sensors.

The AREA spatial filter, unlike the SHARP filter, acts on only the first principal component or the vector selected in a VECT command. It will not act on a NORML display, or on a band selected with a BAND command. The AREA filter is intended to enhance small or narrow features by increasing the brightness contrast between them and their surrounding area (hence the name of the command). Since the hue and saturation information is not modified by the filter, the user can still accentuate edges and boundaries seen by the first principal component without losing any colour information. This filter has two parameters, and it can be used either as a high pass filter (sharpening) or as a low pass filter (blurring). The user changes the filter parameters by typing AREA.

```
!area
AREA DECAY FACTOR (0=<F<1)
```

The user must type a value between 0 and 1. The filter If the user specifies a zero decay factor, the filter is turned off. It is best that the user avoids decay factors larger than 0.95 or less than 0.0 unless he understands what he is doing. Some special effects can be obtained by using unusual decay factors, but they are not ordinarily useful in image interpretation.

WEIGHT OF AREA EFFECT =

The parameter determines the proportion of the difference between the pixel and the surrounding average that is added to the pixel brightness. If the weight is called W, the pixel value P, and the surrounding average AV, then the new pixel value is given by $P_{NEW} = P * (1 + W) - AV * W$. Positive values of W will boost the high spatial frequencies, negative values will tend to attenuate those frequencies, a zero value will effectively turn off the filter. It is recommended that the user stay inside the range of plus and minus 2.

Though the program has no error checking on the decay and weight factors, the user should avoid leaving the suggested ranges unless he knows exactly what he is doing. The filter is quite rudimentary, and does not take care of normalization or ensuring that the output lies within the display range of the MAD (0-63). The filter coefficients are already set at close to their best values (decay factor = 0.8 and weight factor = 0.3) so it is unlikely that the user will improve the enhancement substantially.

The program computes the surrounding average using a recursive filter (i.e. output of the filter is feed back into the input).

$$X'(I) = (VS * X(I-1) + X(I)) / (1 + VS)$$

where VS is the decay factor selected by the user. The filter runs through the scan line in both directions (left and right) and their outputs are combined. The filter is applied again going down across scan lines.

2.3.3. SUPERVISED CLASSIFICATION

!truth

One of the more important applications of digital processing of LANDSAT images is the production of thematic maps. Pixels are assigned to one of several ground cover classes on the basis of their spectral intensities and are then appropriately colour coded. Several different classification schemes have become well established and vary in the degree of interaction between the user and the computer. In the supervised scheme, the user decides which classes he wants to map and must find training samples which define the spectral characteristics of these classes. In order that the user obtain accurate thematic maps, the classes should be distinguishable on the basis of their spectral intensities.

Past experience has shown that the enhancement programs are valuable in choosing the ground cover classes and training samples. The enhancement displays almost all the spectral information available so that the user can tell by just visual examination which classes can be discriminated. It is also easier to select training areas in the enhanced image instead of the normal display. The user therefore has been given the option of running the enhancement programs in training mode using the TRUTH command.

Before typing TRUTH, the user should read the TRUGEN.HLP documentation file, set the SETPAR parameters TRUTH and STATI to the desired file names and display an enhancement preferably without spatial filtering or smoothing (use NOSHARP and AREA DECAY = 0.) The reason it might be preferable to use NOSHARP is that the classification programs do not correct for sensor overlap, and by using the default sharpening parameters, the user might be misled into thinking something was classifiable that would actually be obscured by sensor overlap. It may well be, however, that NOSHARP obscures the user's view of boundaries that might influence the accuracy of the training areas. If the region is such that small features could be important, then both the default SHARP and the default AREA should be used.

Once the user types TRUTH, the graphics cursor is displayed and the program proceeds like TRUGEN.

2.4. PIXSET COMMAND SUMMARY DESCRIPTION

Locating and scaling commands

LINE set the line number at which to begin scrolling the picture. This will be the top line in the picture.

PIXEL set the leftmost pixel of the picture.

MAG indicate the number of screen pixels along a line corresponding to one sensor pixel. The number of lines corresponding to one sensor line will be 1.5 times this number.

GEOM set use of MAG and automatic 1.5 to 1 line to pixel scaling ratio.

NOGEOM turn off automatic scaling of lines and pixels, and permit use of PIXSEL and LINSEL commands

PIXSEL select m out of n pixels for display

LINSEL select m out of n lines for display

Commands to alter the character of the displayed image

ENHANC display a principal components enhancement

XFM select a transformation matrix for the enhancement

NORML display a colour composite of bands 4,5 and 7. Subquestion asked by NORML is whether you want a log, linear or raw display. Log is most common. Linear is scaled from 10 to 63, rather than 0 to 63, and raw comes direct from the sensors without calibration or correction.

VECTOR display one vector in black and white. The selection is as a response to the subsidiary question "VECTOR NUMBER". If two or three vector numbers are given, a (nonstandard) principal components enhancement is displayed.

BAND display one sensor band as a black and white image. NORML should be typed first, to permit setting log, linear or raw transform.

DRAMA change the scale of colour in an enhanced image.

Filtering commands

SHARP Compensate for neighbour pixel overlap

NOSHARP Don't compensate

SMOOTH When magnifying, do linear interpolations between pixels

NOSMOO Don't interpolate, replicate between pixels

AREA High-pass filter by subtracting a regional average from each pixel. Parameters are: exponential decay factor by which each pixel value is multiplied as distance from test pixel increases, and weight applied to average for subtraction. Default values are 0.8, 0.3.

Grid commands

GRID? query current status of grid location and origin
GRID set grid spacing
GRIDO set grid origin
NOGRID remove the grid from subsequent pictures

Miscellaneous commands

WRITE write a display-compatible file for later reference.
Files can be created using this command without actually making a picture on the MAD, by going to the monitor and assigning MAD to the null device (.AS MAD NUL)

TRUTH prepare files for supervised classifier
RESET reset all parameters to default values
LIMIT change limits on spectral bands. Pixels whose value in any of the four bands falls outside the limits will be displayed as grey.

TEST display a coloured test pattern representing the range of colours in the two dimensions. Reading from 12 o'clock, the centres of the sides of the test square should be yellow, green, blue, and red. The middle should be grey.

STATUS list the status of several of the parameters
ROLL start a picture scrolling on the display, and/or writing to a display-compatible file.

EXIT same as control-C. Return to monitor.
HELP list the available commands
NEW list significant changes since date printed on entry to PIXSET.

Grid commands

GRID? query current status of grid location and origin
 GRID set grid spacing
 GRID? set grid origin
 NOGRID remove the grid from subsequent pictures

Miscellaneous commands

WRITE write a display-compatible file for later reference.
 Files can be created using this command without any-
 ally making a picture on the MAD. By going to the moni-
 tor and assigning MAD to the null device (AS MAD NULL)
 prepare files for unprocessed displays
 RESET reset all parameters to default values
 LIMIT change limits on spectral bands. Pixels whose value is
 any of the four bands falls outside the limits will be
 displayed as grey.
 TEST display a coloured test pattern representing the range of
 colours in the two dimensions. Reading from 12 o'clock,
 the centre of the sides of the test square should be
 yellow, green, blue, and red. The middle should be
 grey.

STATUS list the status of several of the parameters
 WFL start a picture scrolling on the display, and/or writing
 to a display-compatible file.
 EXIT same as control-C. Return to monitor.
 HELP list the available commands
 NEW list significant changes since date printed on entry to
 PIXSET.

CHAPTER 3

PROGRAM SOFTWARE

The software is mostly written in FORTRAN, but it calls a number of assembler language routines to perform bit manipulation, byte packing and unpacking, sending data to the MAD, and I/O to disk and tape. The support routines are stored in the libraries, CCRSL/LIB, VIDMOD/LIB, TAPLIB/LIB, and MADMOD/LIB. A programmer considering putting these programs on another system should have an equivalent set of routines, and in particular subroutines to read the JSC LANDSAT data and write it on a colour display.

The software is fairly large, however if one is willing to throw out the many frills such as haze correction, geometric correction, smoothing, and spatial filtering, then it becomes a manageable size. The key programs are EIGENC.FOR, MNVAR.FOR, COLORV.FOR, CSQT.FOR, CSPACE.FOR and some matrix manipulation subroutines. The rest of the programs take care of user dialogue, reading and writing the data. The remaining section describes the functions of the various modules and shows how they interact with each other. The description is very brief and does not reiterate the detailed internal documentation stored in the source files.

It is intended that the following be read in conjunction with the program source listings.

3.1. CCRSL LIBRARY

The following routines are in the CCRSL library as of 18 May 76. Not all of them are used by the routines in the enhancement suite, but the library calls are so embedded in the routines and the other libraries they call that it would be impractical to detail which of the basic CCRSL routines are and which are not used.

BITLOG - RETURNS LOGICAL VALUE OF ONE BIT IN ONE WORD
BITOFF - TURNS OFF ONE BIT IN ONE WORD
BITON - TURNS OFF ONE BIT IN ONE WORD
BITVAL - RETURNS VALUE OF ONE BIT IN ONE WORD
BLT - MOVE BLOCKS OF DATA
BYTPAC - PACK WORDS INTO BYTES
BYTRAN - TRANSFER STRINGS OF BYTES
BYTRLU - TRANSFER STRINGS OF BYTES WITH LOOKUP TABLE TRANSFORM
BYTUPC - UNPACK BYTES INTO WORDS
CHECK - CHECK FOR PRESENCE OF FILES
CLOCK - READ OUT CPU TIME COUNTER
CLOCK1 - RESET CPU TIME COUNTER
CLOSEE - CLOSE CHANNEL OPENED BY INITT/INITT2
DASH - DO ARITHMETIC SHIFT ON DOUBLE-PRECISION WORD
DELETE - DELETE A FILE
DIGINT - INITIALIZE DIGITIZER
DLSH - DO LOGICAL SHIFT ON DOUBLE-PRECISION WORD
DROT - DO ROTATE ON DOUBLE-PRECISION WORD
EXIT1 - SILENT EXIT FROM FORTRAN
GET - DO RANDOM INPUT FROM DISK
GETDIG - GET NEXT DIGITIZER INPUT CHARACTER
IASH - DO ARITHMETIC SHIFT ON ONE WORD
IFLD - TAKE BITS OUT OF A WORD
ILSH - DO LOGICAL SHIFT ON ONE WORD
IMBCCR - BINARY ARRAY CORRELATOR - TYPE 2
IMBCOR - BINARY ARRAY CORRELATOR - TYPE 1
INITT - OPEN A PRE-EXISTING FILE FOR RANDOM INPUT ONLY
INITT2 - OPEN/CREATE A FILE FOR RANDOM INPUT/OUTPUT
INTCHN - RETURN SOFTWARE CHANNEL NO. FROM INIT PACKAGE
INTQCN - QUERY INITT ABOUT SOFTWARE CHANNEL NO.
IROT - DO ROTATE ON ONE WORD
KMPWD - FULL 36-BIT WORD COMPARATOR
OFLD - PUT BITS INTO WORD
PACK - PACK 10A1 INTO 2A5
PUT - DO RANDOM OUTPUT TO DISK
RENAME - CHANGE THE NAME OF A FILE
RUNEX - RUN ANOTHER PROGRAM
SETIBM - SET TAPE TO WRITE IN 'IBM' MODE
SHIFT - MULTI-PURPOSE WORD SHIFTER
SKPEOF - SKIP TO LOGICAL EOF ON TAPE, DECTAPE OR DISK

TLCHN - TRANSLATE FORTRAN LOGICAL UNIT TO SOFTWARE CHANNEL
UNPACK - UNPACK 2A5 INTO 10A1
USETS - DO SUSET. UVO EQUIVALENT OF INIT PACKAGE

3.2. VIDMOD/LIB

This package is used reading the LANDSAT data (in new or old format) from disk and unpacking the data.

DEFVID - defines the name of the VIDEO file, and the structure (disk pack) containing the file.

RLNVID - Reads a line of video data in random access mode. A line contains 3600 pixels (6 or 8 bit quantization), from the 4 Landsat bands, packed into 2409 or 3648 words (6 or 4 pixels per 36-bit word).

DCMVID - Defines line and pixel decimation factors for unpacking routine.

SPCVID - Defines line format, input and output arrays, start pixel, number of pixels for unpacking.

UPKVID - Unpacks the data.

3.3. TAPLIB/LIB

Taplib is a library of subroutines designed for flexible dump mode access to magnetic tapes (read and write). One important reason for the implementation of the routines rather than using the FORTRAN I/O was to allow the programmer better flexibility on error handling. Such things as direct access to the device status word in case of error, and the ability to control bit 29 on output (to determine action on write errors) are available. In addition, the routine "CCTLIN" allows the FORTRAN program to input all 4 tape records corresponding to a particular Landsat scan line (from a JSC Format CCT). The result is placed in a single buffer in core in exactly the CCRS UNIDISK Format. Hence, this may be used for rapid creation of disk files. Also, this format is compatible with one of the available inputs to "UPKVID" or "PAKVID" (in the library VIDMOD). Hence data may be easily re-formatted for processing in FORTRAN programs or repacked for the MAD display.

THE ROUTINES ON THE LIBRARY ARE:

TAPOPN	OPENS CHANNEL FOR INPUT OR OUTPUT
DMPIN	DUMP MODE INPUT
DMPOUT	DUMP MODE OUTPUT
REWIND	REWIND TAPE
UNLOAD	REWIND AND UNLOAD TAPE
BAKREC	BACKSPACE ONE RECORD
BAKFIL	BACKSPACE ONE FILE
WRTEOF	WRITE AN END-OF-FILE
SKPREC	SKIP ONE RECORD
WRT3IN	WRITE 3 INCHES OF BLANK TAPE
SKPFIL	SKIP ONE FILE
SKPEOT	POSITION TAPE AT END-OF-TAPE
CCTLIN	READ 4 RECORDS OF 1 NEW FORMAT
CCRS LANDSAT CCT.	

3.4. MADMOD/LIB

Routines supporting the MAD display

MADOPN - Opens Channel to the MAD

MADLIN - Sends a line of data to the MAD. A line consists of 279 36-bit words. Each word contains the blue, green, and red intensity values for two pixels (3 6-bit bytes per pixel).

SENDA - Sends the ANCILLARY TABLE to the MAD. Used for putting the MAD into colour mode.

GETANC - Gets the ancillary lookup table from a disk file.

ENHANC.FOR

- Main program to ENHANC

- Calls subroutines

TDGO
EXTSET
SUNFIL
CALIB (CALIBN.FOR)
PICDEF
CHECK (CCRS/LIB)
RUNEX (CCRS/LIB)

ENHANC.FOR determines whether to calibrate the data, or go directly to the programs EIGENC or PIXSET. If calibration is required, the program calls SUNFIL to determine the sun angle correction, calls TDGO to open a channel to the LANDSAT data file, calls PICDEF to set the picture bounds, calls CALIB to determine the sensor calibration lookup tables, and then writes the table on a *.CFL file.

EXTSET.FOR

- Subroutine to set the file extension to a 3 letter code name.

parameters

ANAME	file name
EXT	file extension

- Calls subroutines

UNPACK	(CCRSI/LIB)
PACK	(CCRSI/LIB)

SUNFIL.FOR

- Computes sun angle correction factor
- results recorded in SUN.FAC file.

TDGO.FOR

- Reads or starts up reading of LANDSAT video data from a disk file.

entry TDGO

- Starts up reading of LANDSAT video data

- Calls subroutines

INITT (CCSL/LIB)
GET (CCSL/LIB)
CLOSEE (CCSL/LIB)
DEFVID (VIDMOD/LIB)
TAPOP (TAPLIB/LIB)

- Output parameters

DEV - device (T or D) for tape or disk
IND - error return

entry TDLIN

- reads line of LANDSAT data

parameters

LINE required line number
MSS integer array containing 4 bands of packed intensities.
NSN sensor number (1-6)
IND error return flag

- Calls subroutines

RLNVID (VIDMOD/LIB)
DMPIN (TAPLIB/LIB)
BAKREC (TAPLIB/LIB)
SKPREC (TAPLIB/LIB)

PICDEF

- Queries user for first and last lines and pixels to read.

- Output parameters

NLST start line
NEST start pixel
NLEND stop line
NEEND stop pixel

CALIBR.FOR

- produces lookup table for sensor calibration, identification, base correction and sample correction.

- uses histogram matching technique

- input parameters

WSTP start line
WSTL number of spectra (groups of 6 lines)
WSTT start pixel number
WSTL number of pixels

- output parameters

ELUT real array dimensioned 750 containing lookup table
LUT integer array dimensioned 750 containing lookup table
MSS integer array dimensioned at least 700 for holding a line of packed intensities
ISRA error return flag

- call subroutines

TITLE (TITLE FOR)
HSPIL (HSPIL MAC)
SUBCOR

- produces *.GIF files

CALIBN.FOR

- produces lookup table for sensor calibration, linearization, haze correction and sunangle correction.

- uses histogram matching technique

- input parameters

- NSKIP start line
 - NSWT number of swaths (groups of 6 lines)
 - NPXST start pixel number
 - NPXL number of pixels

- output parameters

- FLUT real array dimensioned 1536 containing lookup table
 - LUT integer array dimensioned 1536 containing lookup table
 - MSS integer array dimensioned at least 2409 for holding a line of packed intensities.
 - IERR error return flag

- Calls subroutines

- TDLINE (TDGO.FOR)
 - HSPXL (HSPXL.MAC)
 - SUNCOR

- produces a *.CFL file.

EIGENC.FOR

- Main program to EIGENC

- Calls subroutines

```

      INITT  (CCRSL/LIB)
      GET    (CCRSL/LIB)
      CLOSEE (CCRSL/LIB)
      CHECK  (CCRSL/LIB)
      IFLD   (CCRSL/LIB)
      RUNEX  (CCRSL/LIB)
      LIMSET (LIMBND.FOR)
      EXTSET
      SELLST
      POSVID
      MNVAR
      DCMVID (VIDMOD/LIB)
      DIAG
  
```

EIGENC either estimates the covariances from the LANDSAT data or combines covariances which have been stored in existing *.XFM files. EIGENC calls POSVID to get the user's sampling parameters and MNVAR to compute the covariance statistics. Once the covariance statistics have been determined, the program computes the principal components and records them in an *.XFM file.

MNVAR.FOR

- Estimates the covariances and means of the LANDSAT 4-band data.

- Input parameters

IVERS version number. Must be 1.
NPXL number of pixels to sample per line.
NLIN number of lines to sample
NSEL 41 word array required by SPCVID subroutine.
NFORM format specification parameter required by SPCVID
NSHIFT shift parameter required by SPCVID
ISW switch passed to LIMBND for fast exit
IFIL preprocessing transformation
 0 log and single pixel
 1 log and average 2X2 areas
 10 linear and single pixel
 11 linear and average 2x2 areas

- Output parameters

AMEAN real array dimensioned 4, containing means of 4
 bands.
COVAR real array dimensioned 10, containing symmetric half
 of covariance matrix (C11, C12, C22, C13,...)
MIN
MAX integer arrays dimensioned 4. Remnant of an older
 program.

- Calls subroutines

SPCVID (VIDMOD/LIB)
RLNVID (VIDMOD/LIB)
UPKVID (VIDMOD/LIB)
BLT (CCRSL/LIB)
EXTSET
FRAD (FRAD1.FOR)
HZHI
HZLOG
LIMBND

LIMBND.FOR

- Compares 4 band intensities against limits
- All arguments passed in common blocks /BLIM/ and /BAND/

entry LIMBND

- Reads BANVEC.LIM whenever necessary
- Compares intensities against limits and returns answer in INX.

entry LIMSET

- Sets limits.

entry LIMRES

- Resets limits

- Calls subroutine
CHECK (CCRSI/LIB)

POSVID.FOR

- Get start, stop, lines, pixels, and decimation factors from user.

- Input parameters

IC should be 0 here

- Output parameters

NPXL number of pixels to read

NPXST starting pixel number

NPXDEC pixel decimation factor

NLIN number of lines

NLNST starting line

NLNDEC line decimation factor

MAG magnification factor - always 1.

FRAD1.FOR

- Subroutine to sensor calibrate old format LANDSAT data

- Input parameters

IVERS must be 1

IN1,

IN2,

IN3,

IN4 integer arrays containing the intensities of a line of data in bands 4, 5, 6, and 7.

ND number of pixels per line

LINCD sensor number (1 to 6)

MFILE name of video file (and *.CFL file)

MSTR name of private pack (not used here)

- Output parameters

R1,

R2,

R3,

R4 real arrays containing the sensor corrected data in bands 4, 5, 6, and 7.

IERR error return flag

- When the subroutine is called for the first time, it calls subroutine CGLUT which reads in the sensor calibration lookup table. The intensities and sensor number index the lookup table, and the corrected intensities are recorded in arrays R1, R2, R3, and R4.

HZLOG.FOR

- Subroutine to logarithmically transform the LANDSAT intensities.

- Input parameters

R1,

R2,

R3,

R4 real arrays containing intensities in bands 4, 5, 6, and 7 to be transformed.

ND number of pixels per line.

- Output parameters

RR1,

RR2,

RR3,

RR4 real arrays containing transformed intensities

If called for the first time, the subroutine creates a lookup table with the logarithmic transformation.

HZLIN

- Subroutine to scale the LANDSAT intensities to the range 10-63 for visible image on the MAD. This is the only purpose of the subroutine.

IHZFIX.FOR (IHZFIX and HZFLO)

IHZFIX - subroutine to convert floating pointing intensities to integer.

HZFLO - subroutine to convert integer intensities to floating point.

HZHI.FOR

- Subroutine for performing spatial high-frequency emphasis on data line by line. It assumes that the pixels overlap horizontally and that each contributes a portion to its neighbours' received intensities. The subroutine corrects this by the convolution filter,

$$\text{OUT}(j) = (V \cdot \text{IN}(j) - \text{IN}(j-1) - \text{IN}(j+1)) / (V-2)$$

where $V = 2+6/\text{SV}$ and SV is the sharpening parameter.

- Input parameters

R1,
R2,
R3,
R4 input intensities from a 4 band line.
SV sharpening parameter
ND number of pixels in a line

- Output parameters

RR1,
RR2,
RR3,
RR4 output intensities from the spatial filter.

DIAG.FOR

Subroutine to compute the eigenvalues and eigenvectors of a 4 dimensional matrix

- Parameters

H(4,4) matrix to be diagonalized. On return H contains the eigenvalues along the diagonal, and zeros on the off diagonal.

U(4,4) contains the transformation which diagonalized the original matrix. U(1-4,J) is the Jth eigenvector.

- Calls subroutine JACOBY from the Scientific Subroutine Package.

PIXSET.FOR

- Main program to PIXSET

- Calls subroutines

CHECK (CCRSL/LIB)
INITT (CCRSL/LIB)
GET (CCRSL/LIB)
CLOSEE (CCRSL/LIB)
RUNEX (CCRSL/LIB)
INITT2 (CCRSL/LIB)
EXIT1 (CCRSL/LIB)
DEFVID (CCRSL/LIB)
GETANC (MADMOD/LIB)
MADOPN (MADMOD/LIB)
SENDA (MADMOD/LIB)
EXTSET
ANSWER
LIMSET (LIMBND.FOR)
TRUPIX

(1) PIXSET reads the PX.COM file if it already exists. Otherwise it sets the PX.COM parameters to their default values. (2) PIXSET reads the CONPAR file to get the name of the video file and format specifications. (3) PIXSET initializes the MAD display and (4) types !. (5) On reading the user's command, PIXSET executes the appropriate section of code consisting of operator dialogue. (6) All commands except ROLL and EXIT eventually return back to the ! mode. (7) For the ROLL or EXIT commands, PIXSET rewrites the PX.COM file and goes either to PIXROL or to the monitor.

PIXROL.FOR

- Main program to PIXROL

- Calls subroutines

CHECK (CCRS/LIB)
INITT (CCRS/LIB)
GET (CCRS/LIB)
PUT (CCRS/LIB)
INITT2 (CCRS/LIB)
BLT (CCRS/LIB)
BYTPAC (CCRS/LIB)
RUNEX (CCRS/LIB)
GETANC (MADMOD/LIB)
MADOPN (MADMOD/LIB)
SENDA (MADMOD/LIB)
DEFVID (VIDMOD/LIB)
UPKVID (VIDMOD/LIB)
GEOSET
EXTSET
GETLIN
FRAD (FRAD1.FOR)
HZFLO
HZHI
HZLOG
HZLIN
FEAT
IHZFIX
COLORV
CSQ
INTRPT

(1) PIXROL reads the PX.COM file if it already exists. Otherwise it sets the PX.COM parameters to their starting values. (2) PIXROL reads in the CONPAR file and (3) initializes the MAD. (4) PIXROL reads the *.XFM file and (5) calls GEOSET to set the line and pixel decimation factors. (6) PIXROL reads and unpacks a line of LANDSAT data, (7) preprocesses the line (sensor calibrate, linearize, log transform, spatial filter, principal component transform, colour transform) (8) resamples the line in the horizontal direction to obtain the desired magnification factor (9) resamples the line in the vertical direction to get the desired magnification, and (10) sends the line to the MAD and possibly the Display compatible File. If the user hits any teletype character, PIXROL returns to PIXSET, otherwise it returns to (6) and proceeds with the next line.

GEOSSET

- Subroutine to compute the pixel and line decimation factors in order to get the desired magnification / reduction factor.
- GEOSSET calls DCMVID to pass this information to the VIDMOD package.
- All input and output parameters are passed through the common block /SPEC/.

GETLIN

- Subroutine to get a line of video from disk. It calls RLNVID.

COLORV.FOR

- Subroutine to perform the principal component colour transformation and spatial filtering.

- Input parameters

NVERS must be 5

R1,

R2,

R3,

R4 band 4, 5, 6, and 7 intensity values stored in real arrays.

NPXST first pixel number.

IX sensor number

IVC vector number (0 for colour enhancement, 1,2,3,4 for black and white display of vectors 1,2,3,4.

ISW LIMBND parameter for fast exit.

TNFM name of *.XFM transformation

VS,VW decay and area parameters for spatial filter.

DRAMA colour scale factor

NBIT number of bits in output data.

- Output parameters

IN1,

IN2,

IN3,

IN4 output intensity values for blue, green, and red guns

IERR error return.

COLORV calls subroutine CSQT to generate lookup tables for the colour transformation if not already done. It computes the first principal component of a line of data (or desired vector). It spatial filters the vector. It determines the other principal components and colour transforms them, checking in the mean time whether the intensities satisfy the LIMBND limits.

COLORV has another entry point, CSQ for displaying the colour space using the TEST command.

CSQT.FOR

- Subroutine for creating the colour transformation lookup tables. The lookup tables are stored in 5 arrays in the common block /KTABL/.

LTAB(64) brightness gain factor.

KTABB(1024), KTABG(1024), KTABR(1024) Blue, green, red gun components. The second and third principal components are each converted to a 5 bit number. The two numbers are concatenated into a 10 bit number ranging between 1 and 1024 which is the index into these arrays. The brightness gain factor multiplies the other components. (Integer arithmetic is used for greater speed.)

ISQTAB(1024) a compression table.

- Calls subroutine CSPACE which generates the colour space.

CSPACE.FOR

- Subroutine to compute the red, blue, green values in floating point, given X, and Y values between plus and minus one. In normal usage, the X-Y values will be the values of the second and third principal components of the enhanced data.
- The data space is transformed to match (more or less) the Halsey-Chapanis (1954) identification space.

The routine works in several stages.

(1) It computes useful constants, such as differences of colour space coordinates.

(2) It takes the X,Y data values which range from -1 to +1 and maps it into another square (XXT, YYT) over the same range. The reason for this is so that the original X,Y data map eventually into the red-green and blue-yellow directions, which are corner to corner directions in the colour space.

(3) The square is mapped into the colour triangle defined in the CIE space by the coordinate XR,YR (red phosphor), XG,YG (green), XB,YB (blue), and a bisector of the red-green side XY,YY (yellow). The four corners defined by these four points map the corners of the data square remapped in stage 2. Distances along the four sides map non-linearly, using a cubic transformation approximating the data of Halsey and Chapanis (J. Opt. Soc. Amer, 1954, 442-454) into distances along the triangle sides. XXT (dimension 1) maps onto the purple (red to blue), and the lime (yellow to green) sides while YYT (dimension 2) maps onto cyan (blue to green) and orange (red to yellow). The coordinates XXT, YYT thus define four distinct distances along the four sides of the triangle where the red-green is considered as two sides (red-yellow and yellow-green). The four points thus defined are connected in the CIE space of the triangle by two straight lines (purple to lime, cyan to orange), and their intersection defines the CIE coordinates of the desired colour (AX, AY). These coordinates are available as output parameters from the routine.

(4) The CIE coordinates AX, AY, are converted into proportions of red, green, and blue phosphors required to achieve the desired colour. These are R1A, G1A, and B1A, and they sum to unity.

(5) The MAD screen does not give output intensities linearly related to the numbers provided, so a fudge factor is applied to the phosphor values.

(6) A correction for the different apparent brightness of the different colours is applied in an attempt to ensure that the apparent brightness is the same no matter what colour is produced. This correction is not very good at present, and is not supported by experimental data from the psychophysical literature.